TEST REPORT

Tolly.

#221101

January 2021

Commissioned by

Citrix Systems, Inc.

Citrix ADC CPX

Performance vs. Envoy Proxy

EXECUTIVE SUMMARY

Application Delivery Controllers (ADCs) are the invisible workhorses of distributed computing. Containerized, microservices-based applications can place high demands on ADCs and ADC performance is an essential element of application efficiency.

Citrix ADC CPX (formerly NetScaler) is designed to enhance end-user experience by providing high performance traffic management coupled with low latency in containerized microservices-based environments.

Citrix Systems commissioned Tolly to benchmark the performance of the Citrix ADC CPX solution and compare that with the open source Envoy Proxy. Tests were run in an Amazon Web Services (AWS) environment. Tests focused on measuring latency as an indicator of responsiveness, volume data and transaction rate, and, ultimately, end-user experience. Tests measured P99 latency which measures the latency of the worst 1% of the flows. P99 latency, being the most strenuous test, is the gold standard for latency testing.

The Citrix ADC CPX outperformed the Envoy Proxy in all test scenarios having lower (better) latency and higher data throughput and transaction rates.

THE BOTTOM LINE

Citrix ADC CPX delivers:

1. Only 13% of the latency under stress of Envoy Proxy in data throughput tests using both TLS 1.3 and TLS 1.2 encryption
2. Up to 2.3X data throughput of Envoy Proxy in tests using both TLS 1.3 and TLS 1.2 encryption
3. Up to 3.2X transaction rate of Envoy Proxy in tests using both TLS 1.3 and TLS 1.2 encryption

HTTPS Request Latency with TLS 1.3 & 1.2 for *Volume Traffic Test*

P99 Results

(Lower result is better)

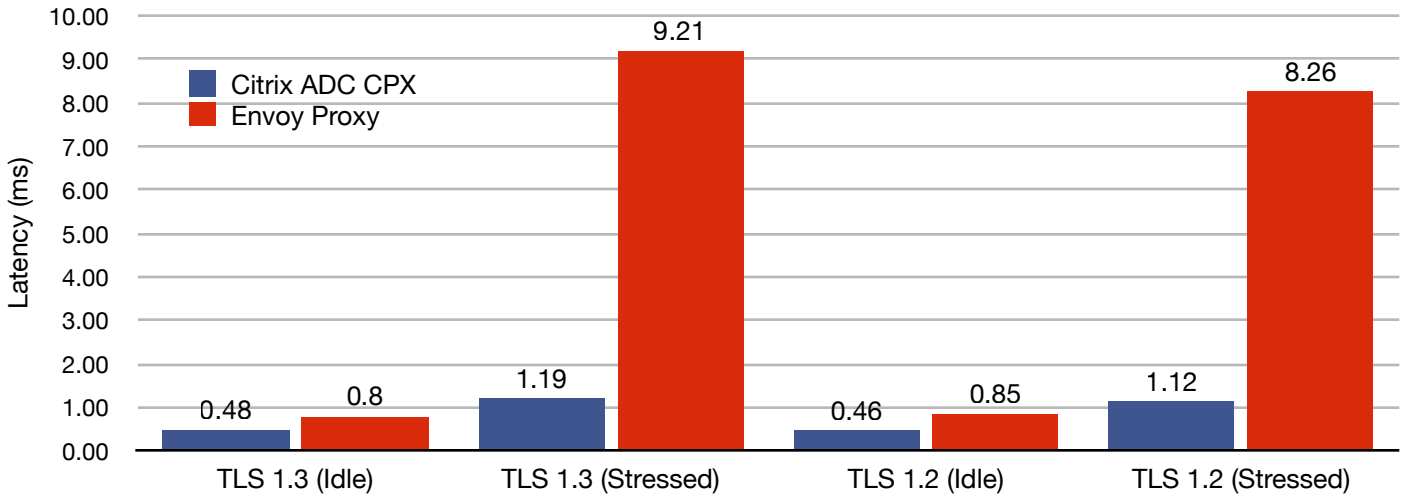

Note: Results as reported by Fortio. Spirent Avalanche Virtual traffic with 9KB response size with CPU for each ADC at ~100% for stressed.

Source: Tolly, December 2020

Figure 1

Key Takeaways

Tests were conducted separately using both TLS 1.3 and TLS 1.2. Systems were benchmarked in "idle" (low traffic) and "stressed" (high traffic) states where CPU use reached 100%.

There are 24 result bars in this report. While the reader is encouraged to review each and every test, the results are consistent throughout. The Citrix ADC CPX delivers better results than the Envoy Proxy in every test.

Citrix ADC CPX latency is better (lower) in every scenario. Lower latency means reduced wait time for distributed applications and, ultimately, end users.

Citrix ADC CPX outperforms Envoy proxy for throughput and transaction rate tests.

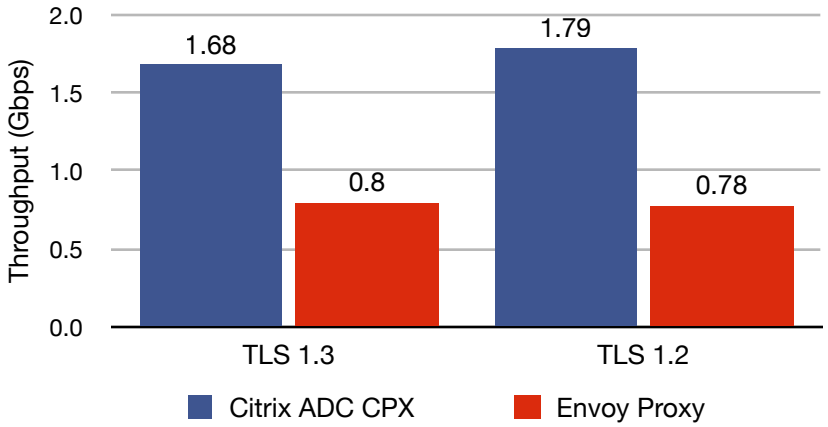

Note: Results as reported by Spirent Avalanche Virtual. ~100% CPU usage.
Source: Tolly, December 2020

Figure 2

Volume Traffic Test

This test measured the latency of the two solutions in an HTTPS volume throughput scenario. Additionally, the maximum throughput was measured. See Figures 1 and 2.

In the volume traffic test, Envoy Proxy latency was 7.7X longer (worse) than Citrix ADC CPX for TLS 1.3 tests and 7.4X longer than Citrix for TLS 1.2 tests under stressed (loaded) conditions.

In terms of data throughput, Citrix ADC CPX delivered 2.1X the throughput of Envoy Proxy in the TLS 1.3 test and Citrix ADC CPX

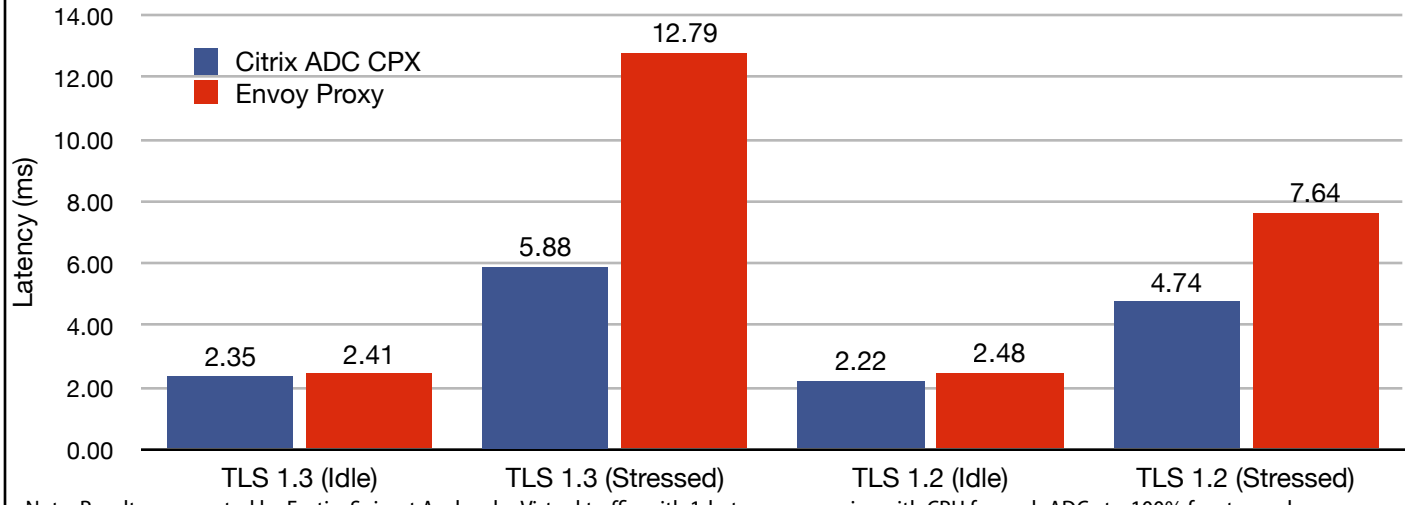

Note: Results as reported by Fortio. Spirent Avalanche Virtual traffic with 1-byte response size with CPU for each ADC at ~100% for stressed.
Source: Tolly, December 2020

Figure 3

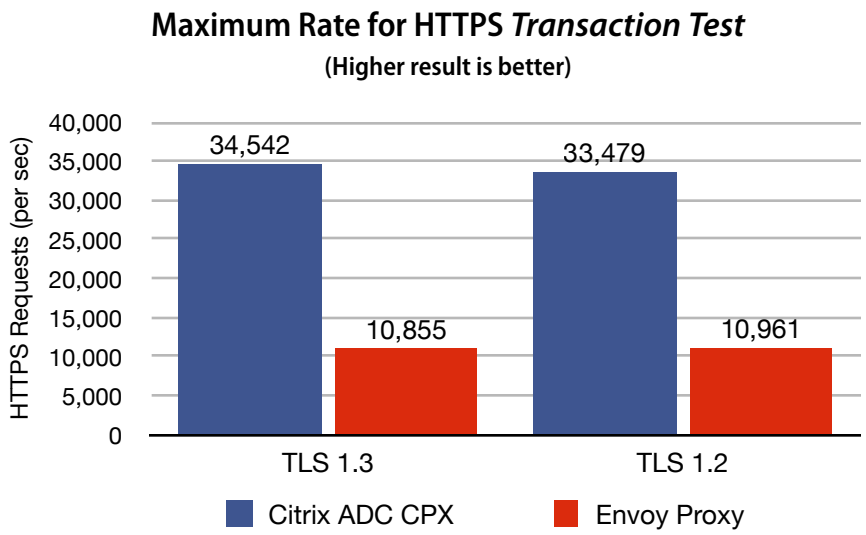

Note: Results as reported by Spirent Avalanche Virtual. ~100% CPU usage.

Source: Tolly, December 2020

Figure 4

throughput was 2.3X that of Envoy Proxy in the TLS 1.2 test.

Transaction Test

This test measured the latency of the two solutions in an HTTPS transaction (small data response) scenario. Additionally, the maximum transaction rate was measured. See Figures 3 and 4.

Envoy Proxy latency was 2.2X longer (worse) than Citrix ADC CPX for TLS 1.3 tests and 1.6X longer than Citrix for TLS 1.2 tests under load.

In terms of transaction rate, Citrix ADC CPX delivered 3.2X the rate of Envoy Proxy in the TLS 1.3 test and Citrix ADC CPX transaction rate was 3X that of Envoy Proxy in the TLS 1.2 test.

Test Setup & Methodology

The focus of the test was to benchmark the performance of containerized application delivery controllers. Benchmarking focused on measuring session latency (delay) in traffic environments designed to model real-world conditions. See Tables 1-4 for details.

Containerized ADCs from Citrix Systems and Envoy Proxy were tested in the Amazon Web Services (AWS) US cloud environment in early December 2020. Tests were run on c5n.xlarge instances. One core was allocated to the ADC (proxy).

ADC Configuration

Since almost all web traffic is now encrypted, all tests were run using encrypted sessions (detailed below).

Citrix Systems, Inc.

Citrix ADC CPX

Containerized ADC Performance

Tolly. Certified

Tested December 2020

P99 Latency

Latency (delay) through the ADC is the primary metric that was used in the test. While some tests report average latency, all latency results in this test are P99 latency. P99 latency is a more stringent measurement that represents the latency of the worst 1% of all flows.

Latency was measured twice in each test scenario, when the ADC was "idle" when processing a light traffic load of 100 requests per second and again under "stressed" conditions when the ADC was handling a heavy traffic load, i.e., CPU use at or near 100%.

Traffic Profiles: Data Throughput & Transaction Tests

It is customary to test both the volumetric (data) throughput and number of HTTPS transactions of ADCs. This allows users to understand the range of performance for a solution. For these tests, the goal was to benchmark the systems at maximum capacity, i.e, at or near 100% CPU usage.

Data Throughput Test

The data throughput test uses a larger response size and, therefore, generates fewer transactions per second than the smaller response size used in the transaction test.

For this test, a 9KB response size was used. Load was increased until the CPU reached 100%. The traffic profile was 100% "GET" requests.

For the "idle" measurement traffic was set to 100 requests (transactions) per second. The "stressed" measurement was taken after the traffic had ramped up to the target rate. Tests ran for at least seven minutes.

Transaction Test

The transaction test uses a smaller response size and, therefore, generates more transactions per second than the larger response used in the throughput test.

For this test, a GET request generated 1 byte HTTPS response. Load was increased until the CPU reached 100%. The traffic profile was 100% "GET" requests.

For the "idle" measurement traffic was set to 100 requests (transactions) per second. The "stressed" measurement was taken after the traffic had ramped up to the target rate. Tests ran for at least seven minutes.

Test Details

Solutions Under Test

Vendor	ADC	Version
Citrix Systems	Citrix ADC CPX	13.0-70.7
Open Source	Envoy Proxy	1.14.2

Table 1

Operating Environment

Component	Name	Version
Operating System	CentOS Linux	7.7 Linux, build 3.10.0-1062.12.1.el7.x86_64
Container Platform	Docker	1.13.1, build 64e9980/1.13.1

Table 2

Test Variables

ADC Configuration	Encryption Protocol - Transport Layer Security (TLS)	
ADC Only	v1.2 TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	v1.3 TLS_AES_256_GCM_SHA384

Table 3

Test Tools

Vendor	Solution	Function
Spirent Comm.	Avalanche Virtual	Traffic generator
Fortio (Open Source)	Fortio	Measure latency

Table 4

Source: Tolly, December 2020

Encryption: Transport Layer Security 1.2 & 1.3

Inter-microservices communications is increasingly encrypted for security purposes. Thus, it is important to benchmark the performance with encrypted (i.e., HTTPS) traffic in order to get an accurate reading on how an ADC/Proxy will perform in the real-world. Thus, all the benchmarking performed used traffic encrypted using TLS (which is the follow-on to SSL).

Today, traffic will consist of sessions using the older TLS 1.2 protocol along with a growing number of sessions using the current TLS 1.3 protocol (which will ultimately replace TLS 1.2). Thus, separate tests were run with each to benchmark encryption protocol specific performance.

Test Procedure

Spirent Avalanche Virtual was used to generate all traffic for the transaction/data throughput tests. Avalanche Virtual was resident in the same datacenter as the ADCs under test and simulates both the client and server sides of the connections.

Fortio was used to measure P99 latency and ran in tandem with Avalanche Virtual.

For the idle measurement, Avalanche Virtual generated 100 requests per second for several minutes. During this time, Fortio was also run and reported P99 latency.

For the stressed measurement, Avalanche Virtual was used, first, to generate a load level that would bring the two vendors SUTs to equivalent load levels. Engineers would then reduce the load by approximately 5%. This was done to be sure that Fortio could run and still stay within the maximum levels established for the test. Fortio was run and reported P99 latency.

Test Systems Summary

Vendor	Product	Web	
Fortio	**Fortio**	https://fortio.org	
Spirent Comm.	**Avalanche Virtual**	https://www.spirent.com	

About Tolly

The Tolly Group companies have been delivering world-class IT services for more than 30 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

You can reach the company by E-mail at sales@tolly.com, or by telephone at +1 561.391.5610.

Visit Tolly on the Internet at:
http://www.tolly.com

Terms of Usage

dj-m-wt-2021-01-04-VerK